

AMES REPORT
112-51-212-
161501
208.

Performance Analysis of FDDI

Marjory J. Johnson

April, 1988

Research Institute for Advanced Computer Science
NASA Ames Research Center

RIACS TR 88.11

(NASA-CR-183206) PERFORMANCE ANALYSIS OF
FDDI (NASA) 20 F CSCI 09B

N88-30321

Unclas
G3/61 0161501

RIACS

Research Institute for Advanced Computer Science

Performance Analysis of FDDI

Marjory J. Johnson

April, 1988

Research Institute for Advanced Computer Science
NASA Ames Research Center

The Fiber Distributed Data Interface (FDDI) is an emerging ANSI and ISO standard for a 100 megabit-per-second fiber-optic token ring. The purpose of this paper is to analyze performance of the FDDI media-access-control protocol using a simulation developed at NASA Ames Research Center. This study includes both analyses using standard measures of performance (including average delay for asynchronous traffic, channel utilization, and transmission-queue length) and analyses of characteristics of ring behavior which can be attributed to constraints imposed by the timed-token protocol on token-holding time (including bounded token-rotation time, support for synchronous traffic, and fairness of channel access for nodes transmitting asynchronous traffic).

RESEARCH REPORT
NASA Ames Research Center
(U.S. GOVERNMENT WORK)
DISTRIBUTION STATEMENTS
UNCLASSIFIED

COMPUTER PROGRAM
DATA BASE, ANALYSIS, AND DESIGN
UNCLASSIFIED
EXCLUDED FROM ABSTRACT

This work was supported by the National Aeronautics and Space Administration under NASA Contract NCC2-387 to the Universities Space Research Association (USRA).

PERFORMANCE ANALYSIS OF FDDI

*Marjory J. Johnson
Research Institute for Advanced Computer Science
NASA Ames Research Center
Moffett Field, California 94035*

1. Introduction

The Fiber Distributed Data Interface (FDDI) is an emerging American National Standards Institute (ANSI) and International Standards Organization (ISO) standard for a 100 megabit-per-second fiber-optic token ring. Performance capabilities of FDDI far exceed those of the IEEE 802.X set of network protocols, thus creating new and exciting possibilities for future networking. NASA has several potential uses for FDDI, including the Data Management System on board the Space Station.

The purpose of this paper is to analyze performance of the FDDI media-access-control protocol. The results contained herein were obtained by using a simulator developed at NASA Ames Research Center as a tool for Space Station developers.* Our analysis consists of two parts. First we analyze FDDI in terms of standard performance measures, including delay, channel utilization, and transmission-queue length. Then we analyze characteristics of ring behavior that stem from constraints imposed by the timed-token protocol on token-holding time.

*This simulator, called LANES (Local Area Network Simulator), is a parameter-driven simulator of the FDDI media-access-control protocol. LANES was developed by the Data Networks Concepts group at NASA Ames Research Center, under the direction of Terry Grant.

Several performance studies of FDDI have been reported in the literature. Ulm [9] was the first to analyze performance of the timed-token protocol. He examined utilization as a function of various ring parameters. Papers by Sevcik and Johnson [8], by Johnson [4,5,6], and by Dykeman and Bux [2] are primarily analytical studies of properties of timed-token behavior. A simulation study by Dykeman and Bux [1] focuses on performance characteristics of individual asynchronous priority classes.

2. FDDI Access Protocol

FDDI is a timed-token-rotation protocol; timers within each node cooperatively attempt to maintain a specified token-rotation time by using the observed network load to regulate the amount of time that an individual node may transmit. Consequently, token-rotation time for FDDI is bounded. This bound is a function of a ring parameter, called T_Opr , which specifies the expected token-rotation time. Because token-rotation time is bounded, FDDI is able to support applications that have stringent requirements on frequency of channel access, such as packet voice and real-time control. FDDI offers two classes of service: synchronous service for applications with stringent channel-access requirements such as those designated above, and asynchronous service for applications which do not have such stringent channel-access requirements. Although the protocol supports multiple priorities for asynchronous traffic, in the analysis reported herein we assume that all asynchronous traffic has the same priority.

For a detailed discussion of the way the FDDI timers work to control access to the channel, see [3,5,7]. Synchronous service is discussed in detail in Section 3.2.2.

3. Simulation Results

3.1. Standard Measures of Performance

The ring configuration used to obtain the results in this section is presented in Table 1. The network is homogeneous; all traffic is asynchronous and each node generates frames at the same specified mean arrival rate. Interarrival times for frames at each node are exponentially distributed.

3.1.1. Average Frame Delay

Figure 1 is a graph of average frame delay versus offered load. The delay measured in the simulation is the time from generation of the frame at the source node to receipt of the frame at the destination node. This includes queueing delay at the source node, transmission time, and the time required for the frame to propagate from the source node to the destination node.* As a lower bound for delay, transmission time for a 4040-byte frame is 323.2 microseconds, and propagation delay is a maximum of 15 microseconds (including a 60-bit internal delay for each node).

Table 1. Ring Configuration

Parameter	Value
Number of Nodes	20
Distance between Nodes	30 meters
T_{Opr}	40 milliseconds
Frame Size	4000 bytes
Header Size	40 bytes

*Delays due to processing at the upper layers of the OSI (Open Systems Interconnection) network model, either at the source or at the destination, are not included.

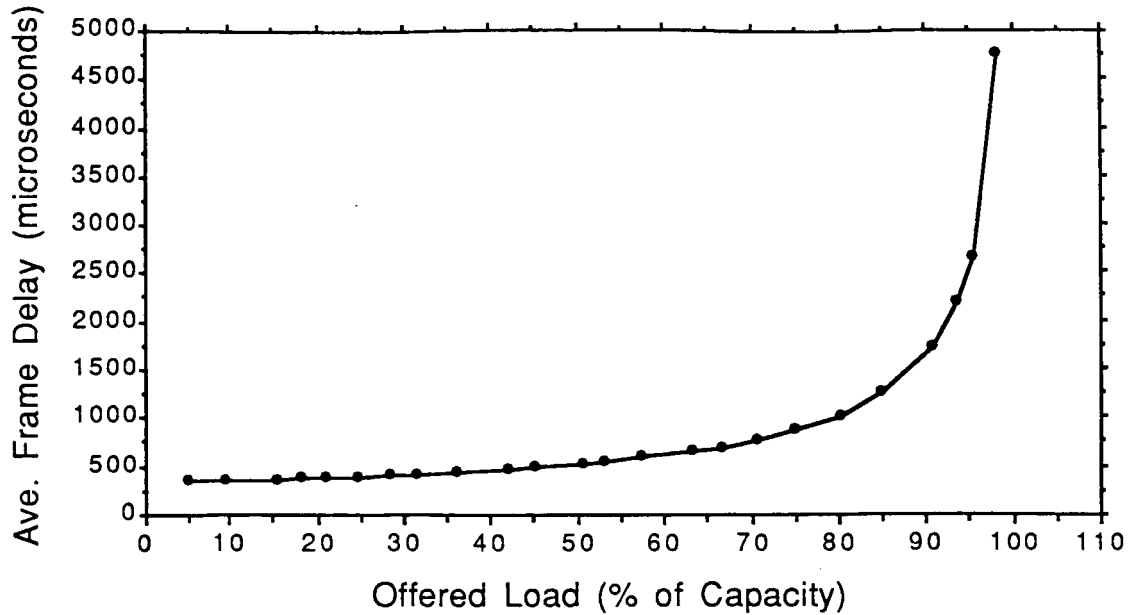


Figure 1. Average Frame Delay vs. Offered Load

Average frame delay increases slowly as a function of offered load. The slope of the curve doesn't increase rapidly until the offered load exceeds 85% of capacity. Even when the offered load is 98% of capacity, the average delay is approximately fifteen frame-transmission times. Maximum frame delays experienced in these runs range from 540.6 microseconds at 5% offered load to 21647.5 microseconds (approximately half the T_Opr value) at 98% offered load.

3.1.2. Channel Utilization

Figure 2 presents utilization of the channel as a function of offered load. Utilization increases linearly until the network is saturated, and then levels off at approximately 99.9%. Ulm [9] presents a formula for ring utilization as a function of ring latency and average token-rotation time. Our figures for utilization at the various offered loads (the associated average token-rotation times are presented in Figure 4, Section 3.2.1) agree almost precisely with this formula.

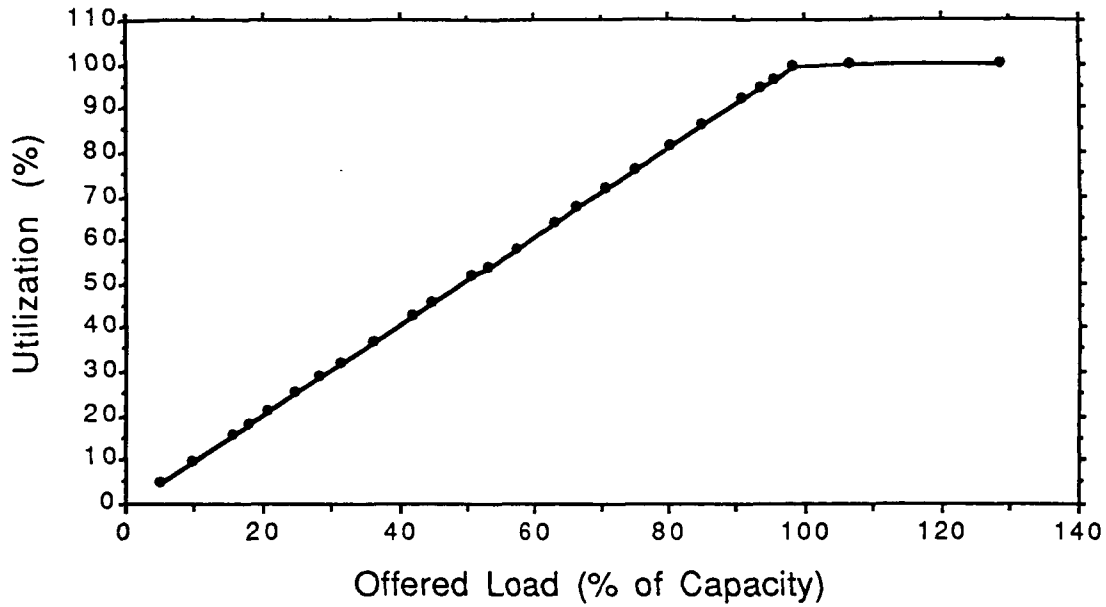


Figure 2. Utilization vs. Offered Load

3.1.3. Queue Lengths

At each node frames are placed into the transmission queue as soon as they are generated, and they remain there until they are transmitted on the channel. Figure 3 plots both average and maximum queue lengths, as functions of the offered load. Since our network configuration is homogeneous, for any given offered load, the average number of frames in the transmission queues at the individual nodes are all approximately the same. Hence, a single value is given in Figure 3, rather than presenting a value for each node. The maximum value given is the maximum over all the nodes.

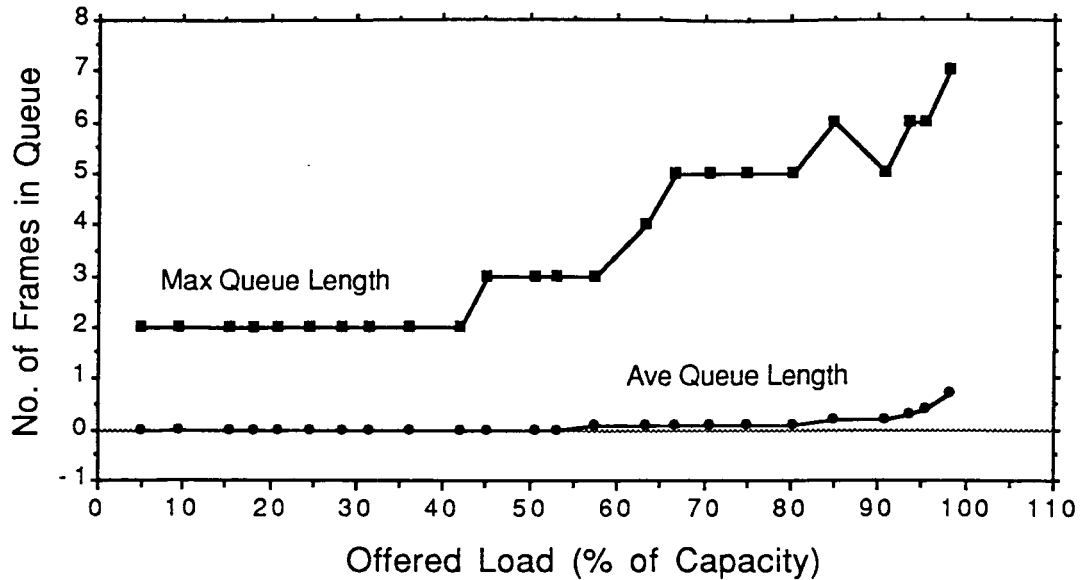


Figure 3. Queue Length vs. Offered Load

The average queue length is 0.0, until the offered load exceeds 50%. The average queue length then rises very slowly, until at 98% offered load, it is 0.7. Maximum queue length is 2, until the offered load exceeds 40%. Even when the offered load is 98%, maximum queue length is only 7. Such low queue lengths, along with the low average delay figures presented in Figure 1, suggest that even when the offered load is as high as 98% of capacity, the ring is able to service all the traffic satisfactorily.

FDDI timers allow for transmission of multiple frames during a single token capture if the preceding token cycle was sufficiently short. Even at 5% offered load, we see a few instances of multiple frames being transmitted during a single token capture. As the offered load increases, the frequency with which multiple frames are transmitted also increases. When the ring is saturated, an average of six frames are transmitted per token capture for our ring configuration. Both

Figure 3 and the above information suggest that performance would be degraded if multiple frames could not be queued for transmission.

3.2. Distinctive Features of Timed-Token Ring Behavior

Distinctive features of ring behavior that stem from constraints of the timed-token protocol on the time that individual nodes can hold the token include a bounded token-rotation time, the ability to provide support for synchronous applications, and the ability to provide equal access to individual nodes for asynchronous transmission. In this section we present simulation data to support earlier analytic results [4,5,6,8].

3.2.1. Bounded Token-Rotation Time

It can be proved that the maximum token-rotation time for any ring configuration is $2 \times T_Opr$ [5,8], while average token-rotation time is less than or equal to T_Opr [8]. Figures 4.a and 4.b are graphs of average token-rotation time as a function of offered load, using the ring configuration of Table 1 with T_Opr equal to 40 milliseconds. The time required for the token to rotate around an empty ring is 15 microseconds. It is only when the offered load exceeds the capacity of the ring that the average token-rotation time approaches T_Opr as a bound. This indicates that the T_Opr value is not a limiting factor for our ring configuration except during bursts of activity that temporarily saturate the network. Figure 4.a clearly illustrates the asymptotic behavior of average token-rotation time as a function of offered load. The relatively flat portion of the graph (when the offered load is below saturation) is enlarged in Figure 4.b, so that the gradual increase in average token-rotation time can be seen.

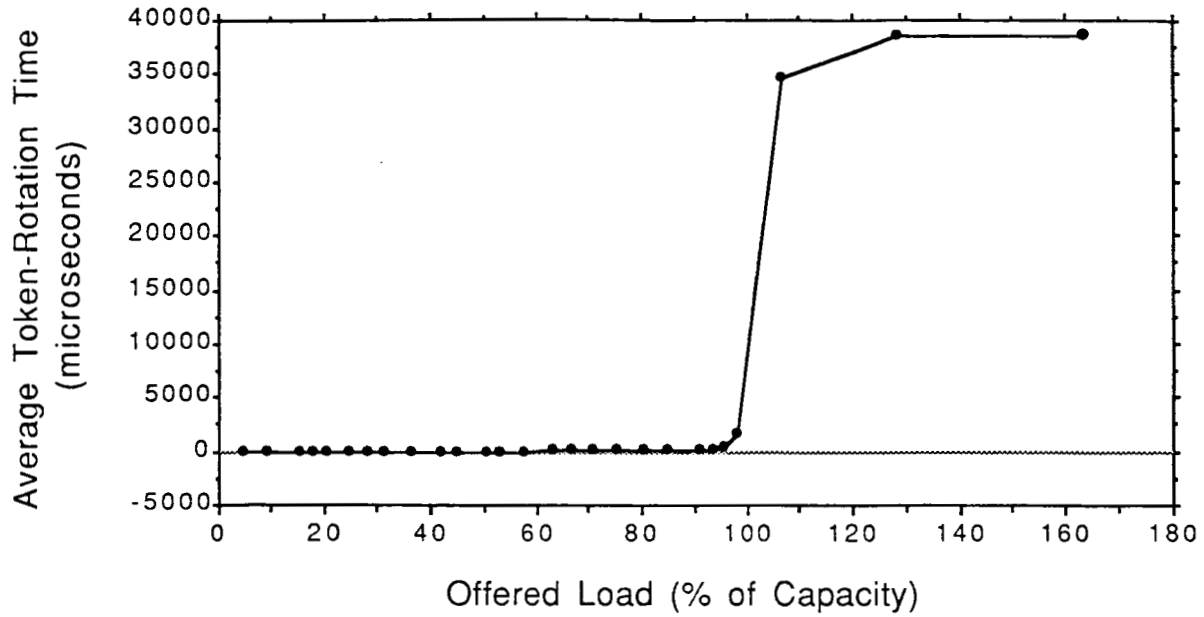


Fig. 4.a. Complete graph, showing asymptotic behavior

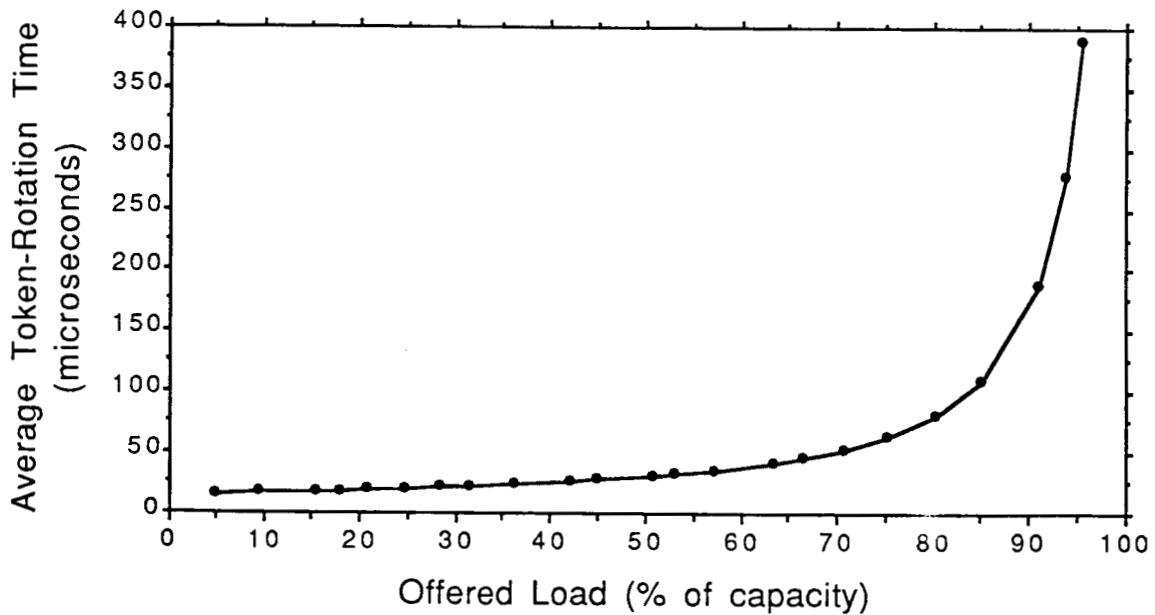


Fig. 4.b. Enlargement of bottom portion of Fig. 4.a

Figure 4. Average Token-Rotation Time vs. Offered Load

Related information that sheds light on the responsiveness of FDDI is the amount of time a node must wait to be serviced when it has one or more frames queued for transmission. Table 2 presents both average and maximum values for a range of offered loads.

3.2.2. Synchronous Service

The FDDI media-access-control (MAC) protocol guarantees a bounded service interval for synchronous applications. That is, at any given time during ring operation, a node which has been assigned some synchronous bandwidth is guaranteed to be serviced within a specified time interval, which we call the *synchronous-access time interval*. This guarantee is possible only because token-rotation time is bounded. At ring initialization, nodes negotiate the value of T_Opr , a ring parameter which specifies the expected token-rotation time, so that synchronous channel-access requirements of all the nodes will be satisfied. The smallest requested value is assigned to T_Opr . Then each node is assigned

Table 2. Wait for Usable Token

Offered Load (% of capacity)	Average Wait (microseconds)	Maximum Wait (microseconds)
10	30	509
20	47	968
30	76	1328
40	133	1913
50	151	2189
60	309	5723
70	421	4800
80	650	6097
90	1367	9469
95	2348	13244
105	30246	38493

an amount of time, called its *synchronous-bandwidth allocation*, for synchronous transmission each time it receives the token. Each time a node receives the token, it may transmit synchronous frames for its allotted time. In contrast, asynchronous transmission is allowed only if the load on the ring is light enough to support it.

The total of all synchronous assignments must not exceed 100 percent of T_Opr , since this is total capacity of the ring. Hence, the value of T_Opr determines the maximum possible volume of synchronous traffic that can be supported in a particular ring configuration. For this reason, it is worthwhile to determine the largest possible value of T_Opr that will provide satisfactory channel access to support the network's synchronous traffic. An additional reason for assigning the largest possible value to T_Opr is that ring operation will be more efficient, although this factor was negligible in experiments conducted for this study. Since the average token-rotation time approaches T_Opr as the offered load nears capacity, clearly T_Opr must be assigned a value less than or equal to the length of the synchronous-access time interval, or synchronous-frame delay could be excessive on a regular basis.

Since the interval between successive token visits to a node is bounded above by $2 \times T_Opr$ [5,8], the standards document [3] states that T_Opr should be set to one-half the desired synchronous-access time interval. In [6] it is shown that maximum token-rotation time for a particular ring configuration is actually dependent on total synchronous-bandwidth allocation for that ring and, consequently, is generally less than $2 \times T_Opr$. A formula is presented therein for computing the largest possible value of T_Opr which will guarantee the desired frequency of channel access for a particular configuration, based on that

configuration's total synchronous-bandwidth requirements. However, since the occurrence of the exact set of circumstances which would cause token-rotation time to assume (or even approach) the maximum value is extremely unlikely, it seems that even this setting for T_Opr may be unnecessarily restrictive. Simulation results are presented in [6] which suggest that for some synchronous applications, sufficient support is provided when T_Opr is set equal to the length of the desired synchronous-access time interval. We present similar results here, using the ring configuration described in Table 3.

In this configuration fifteen of the nodes are synchronous nodes, i.e., they generate only synchronous traffic, while the remaining five nodes are asynchronous nodes, i.e., they generate only asynchronous traffic. Each synchronous node generates synchronous frames at a constant rate of one frame every 6750 microseconds. While the interval between consecutive synchronous frames generated at any individual node is constant, the generation of synchronous frames at different synchronous nodes is staggered. For each synchronous node, it is desired that any given frame at that node should be transmitted before the next one at that same node is queued for transmission, i.e., the desired synchronous-

Table 3. Ring Configuration

Parameter	Value
Number of Sync Nodes	15
Number of Async Nodes	5
Interarrival Time between Sync Frames	6750 μs
Distance between Nodes	30 meters
T_Opr	6750 μs
Length of Synchronous- Access Time Interval	6750 μs
Sync Bandwidth Allocation	75%

access time interval is 6750 microseconds. Each synchronous node is allocated synchronous bandwidth to transmit exactly one synchronous frame each time it receives the token. Thus, total synchronous-bandwidth allocation for the entire network is approximately 75% of the specified T_Opr value. Note that it is theoretically impossible to set T_Opr according to the standards document (or according to the formula given in [6] either), because there is too much demand for synchronous bandwidth relative to the desired frequency of channel access.

We conducted a series of simulation runs using this basic configuration, varying the overall offered load by varying the load offered by the asynchronous nodes. Even with the total offered load as high as 95% of capacity, all synchronous-frame delays were less than 6750 microseconds, i.e., they were all within the desired range.*

When the asynchronous load was increased so that the total offered load was approximately 120% of capacity, 72 out of a total of 2225 synchronous frames (approximately 3.2%) experienced delays that exceeded 6750 microseconds. The most instances of excessive delay occurred for node 12, with 13 of 148 frames, or 8.8%, experiencing delays greater than 6750 microseconds. Figure 5 is a histogram of synchronous-frame delays for this node.

*While there is not a direct correlation between synchronous-frame delay and the time between consecutive channel accesses, as long as synchronous delay does not exceed T_Opr , then the ring is performing as desired.

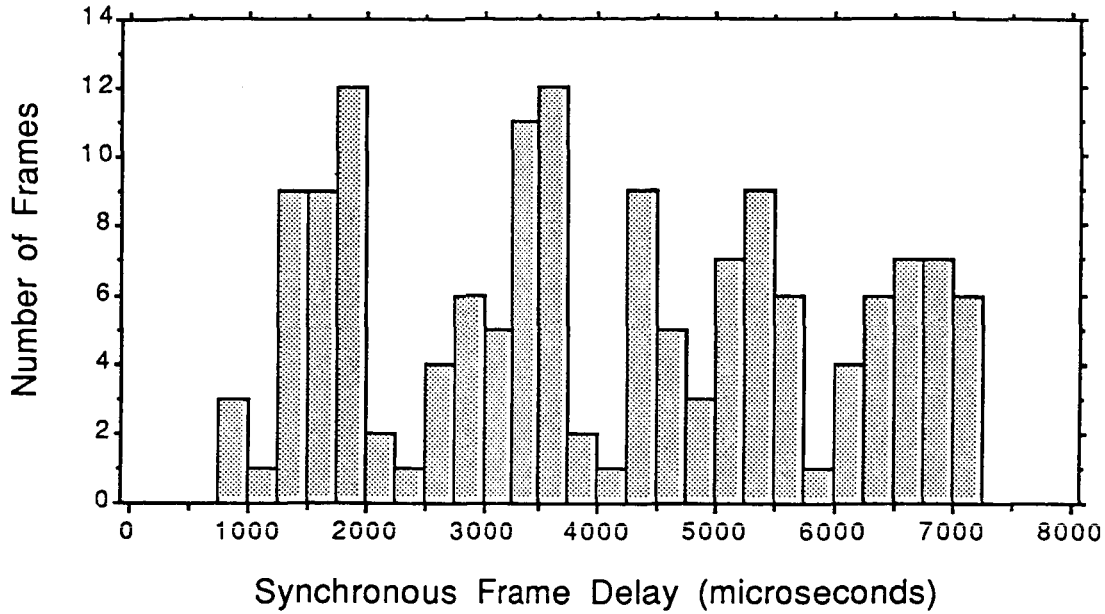


Figure 5. Frequency Distribution of Synchronous Frame Delays for Node 12

As was discussed in [6], excessive delays (i.e., delays greater than 6750 microseconds) tend to occur in clusters, because excessive delay for one frame will cause frames to back up in the transmission queue. Since the synchronous allocation for this particular configuration allows only one frame to be transmitted during each token rotation, since the token-rotation time in a saturated ring approaches T_{Opr} , and since an additional frame is added to the queue every T_{Opr} microseconds, it may take several rotations before the queue becomes empty again. There were five such clusters for node 12 in the above scenario. These clusters of excessive delays can be eliminated by purging a synchronous frame which is pending transmission when a new synchronous frame becomes queued for transmission at the same node. If this purging technique had been implemented in the above scenario, then only five synchronous frames from node 12 (approximately 3.4%) would have been lost, and the rest would have experi-

enced delay within the desired range. Moreover, less than 3.5% of the synchronous frames generated by any single node would have been lost, and 98% of all synchronous-frame delays for the entire network would have fallen within the desired range. Depending on the particular application, this might be entirely satisfactory service for synchronous traffic. It is certainly more desirable than not being able to support the volume of synchronous traffic at all.

3.2.3. Fairness of Access for Asynchronous Traffic

While bandwidth is guaranteed for synchronous traffic, asynchronous traffic is transmitted only if the load on the ring is light enough to support it. According to the standards document [3], the FDDI MAC protocol "supports fair access at a frame granularity" for asynchronous transmission. This would be a desirable property, for it would ensure that individual nodes wouldn't suffer starvation, even during periods of ring saturation. Although the claim in [3] is made without any justification, a proof is presented in [4] that, under some simplifying assumptions, all nodes do have equal access to the channel to transmit asynchronous frames.

Simulation data to support the theory of fairness of channel access is most easily obtained by constructing scenarios in which the offered load exceeds the capacity of the ring. Figure 6 is a histogram of the number of frames transmitted by each of the nodes in a twenty-node homogeneous ring where all traffic is asynchronous and the offered load at each of the nodes is essentially infinite (i.e., the transmission queues always contain frames waiting for transmission). Statistics collected over ten seconds of ring operation show that the largest number of frames transmitted by a single node was 1557, while the smallest number of frames transmitted by a single node was 1530, 98.3% of the larger number. Ring

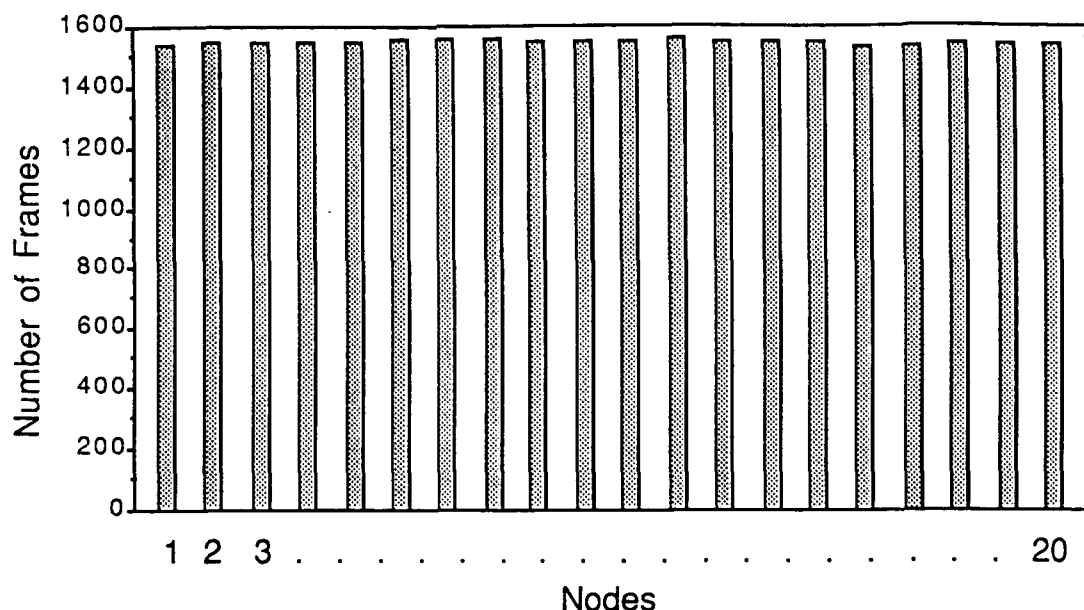


Figure 6. Number of Asynchronous Frames Transmitted by Individual Nodes (homogeneous ring)

operation in this scenario is essentially time-division multiple access (TDMA), with a six-frame time slot for each node during each token rotation. This represents, of course, the most efficient utilization of the channel in a saturated ring. We obtained similar results when we repeated the experiment with some synchronous traffic on the ring.

It is easy to construct scenarios similar to the above which demonstrate equal access for asynchronous transmission. However, there is also simulation evidence to suggest that the pattern of channel access for asynchronous transmission is sensitive to settings of ring parameters and may not always be fair. We ran a simulation with a seven-node ring, in which Nodes 1, 2, and 4, called Class A nodes, transmit frames that are 10% longer than the frames transmitted by the remaining nodes, called Class B nodes. As in the previous scenario, the offered load at each node is infinite. Under conditions of equal

access to the channel, we would expect Class A nodes to transmit X frames each and Class B nodes to transmit Y frames each, where $Y > X$, since the frames transmitted by Class B nodes are smaller. Figure 7 shows that this is not at all what happened. In further experiments, we observed that even a ten-byte difference in frame size for different nodes (a difference which could be caused by some nodes transmitting more idle symbols in the frame header than others) can cause a significant imbalance in channel access.

Dykeman and Bux [1] have detected similar problems of unfairness in scenarios involving multiple priorities of asynchronous traffic. In their simulation experiments they noted that nodes transmitting low-priority frames which are located immediately downstream from nodes transmitting high-priority frames have an unfair advantage over other low-priority nodes.

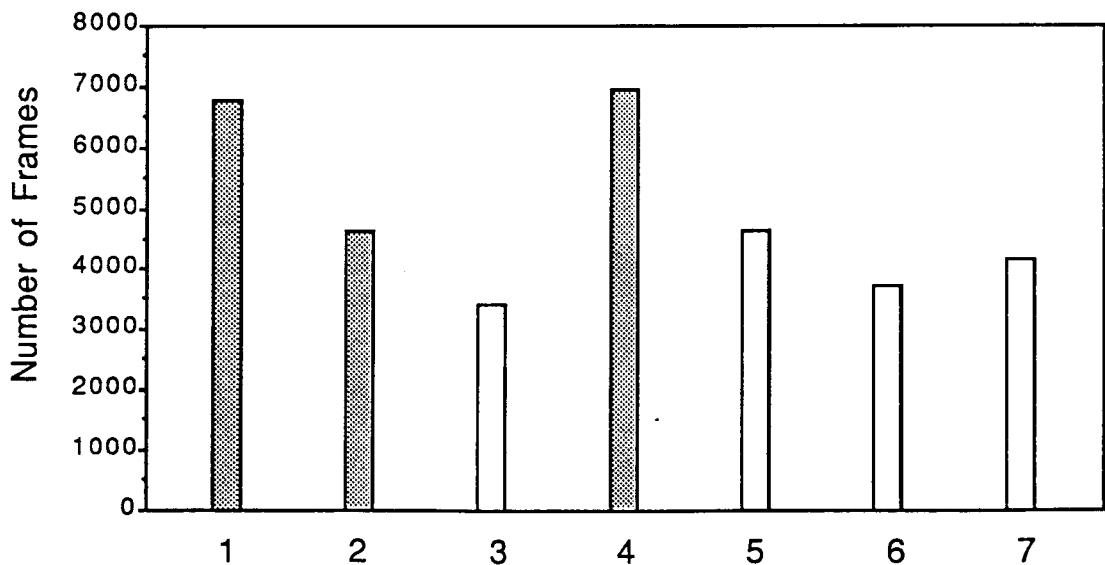


Figure 7. Number of Asynchronous Frames Transmitted by Individual Nodes (different frame sizes)

4. Conclusions

In this paper we have presented a performance analysis of the FDDI token-ring protocol. Using standard measures of performance, we have shown that average frame delay is low until the ring nears saturation; utilization follows the ideal curve, increasing linearly until the ring reaches saturation and then leveling off; and transmission-queue lengths remain small, until the ring reaches saturation, indicating that frames are transmitted almost as soon as they are generated.

In addition, we presented simulation data to support results from earlier analytic studies of distinctive features of the FDDI timed-token protocol. First we demonstrated the asymptotic behavior of average token-rotation time as a function of offered load. Then we investigated synchronous-frame delays with a more relaxed setting of T_Opr than that specified in the standards document. Our results indicate that the service provided when T_Opr is set equal to the desired synchronous-access time interval would be satisfactory for applications which can tolerate excessive delays for a small percentage of synchronous frames. Finally, we presented simulation data to demonstrate that the pattern of channel access for nodes transmitting asynchronous frames is sensitive to values of ring parameters, and that access may or may not be fair, depending on the particular configuration.

References

1. D. Dykeman and W. Bux, "An Investigation of the FDDI Media Access Control Protocol," in *Proc. EFOC/LAN '87*, Basel, Switzerland, (IGI Europe), Information Gatekeepers, Inc., Boston, June, 1987, pp. 229 ff.
2. D. Dykeman and W. Bux, "Analysis and Tuning of the FDDI Media-Access Control Protocol," to appear in *IEEE J. Sel. Areas in Commun.*, vol. SAC-6, July 1988.
3. *FDDI Token Ring Media Access Control*, American National Standard, X3.139 - 1987.
4. M. J. Johnson, "Fairness of Channel Access for Non-Time-Critical Traffic Using the FDDI Token Ring Protocol," in *Proc. Advanced Seminar on Real-Time Local Area Networks*, Bandol, France, INRIA, Apr. 1986, pp. 145-157.
5. M. J. Johnson, "Proof that Timing Requirements of the FDDI Token Ring Protocol are Satisfied," *IEEE Trans. Commun.*, vol. COM-35, no. 6, pp. 620-625, 1987.
6. M. J. Johnson, "Analysis of FDDI Synchronous Traffic Delays," in *Proceedings of Systems Design and Networks Conference*, IEEE, Apr. 1988.
7. F. E. Ross, "FDDI - A Tutorial," *IEEE Commun. Mag.*, vol. 24, no. 5, pp. 10-17, 1986.
8. K. C. Sevcik and M. J. Johnson, "Cycle Time Properties of the FDDI Token Ring Protocol," *IEEE Trans. Software Eng.*, vol. SE-13, no. 3, pp. 376-385, 1987.
9. J. M. Ulm, "A Timed Token Ring Local Area Network and its Performance Characteristics," in *Proc. 7th Conf. Local Comput. Networks*, IEEE, Feb. 1982, pp. 50-56.